REMARKS

Claims 1-4, 6-9, 11-13, 16-17, 19-32, 34-35, 38, 40-41, 44, 46-47, 50, 52-53, 56, and 58-64 are now pending in the application. Claims 5, 14-15, 33, 36-37, 42-43, 48-49, and 54-55 have been canceled without prejudice or disclaimer. Claims 1, 4, 20-22, 46-47, and 50 have been amended, and new claims 61-64 added, without introduction of new matter. Favorable reconsideration is respectfully requested in view of the above amendments and the following remarks.

Claims 1-9, 11, 17, 19, 20-22, 32, 34, 38, 40, 44, 46, 50, 52, 56, 59, and 60 again stand rejected under 35 U.S.C. §103(a) as allegedly being unpatentable over Childers et al. (US Patent 5,986,913) (henceforth "Childers") in view of Campbell et al. (US Patent 5,021,947) (henceforth "Campbell"). This rejection is respectfully traversed.

Claim 1 has been amended to state specifically that the data processing architecture is operable to process more than one data packet at a time. The claim now also spells out more precisely the dynamic distribution feature discussed in earlier responses. In particular, claim 1 defines that "the input device is operable to distribute whole data packets of unpredictable size across one or more of said processing elements." Claim 1 further now defines "a data packet greater than a predetermined size being divided into portions and each portion distributed to a respective processing element; and a data packet less than a predetermined size being distributed to a single processing element; wherein the data processing architecture is operable to process at least one data packet at a time." No new matter has been introduced. Support for these features is set forth in the following discussion, which also explains why the prior art of record fails to disclose or suggest at least these features.

As has previously been discussed in previously-filed responses, Childers discloses a high-speed sense amplifier/memory configuration resistant to voltage spikes. The practical embodiment is a serial video processor in which lines of video are encoded to create packets representing respective pixels of a video line. The processor is a SIMD/RISC device consisting of a one-dimensional array of 1-bit PEs in which there are as many PEs as there are pixels in the video line. Data for each pixel is handled by a respective PE. The packets are all of the same size and it appears inevitable that the allocation of data to the PEs is based purely on *numbers* of pixels and *numbers* of PEs. The Childers processor has clearly been engineered so as to match the number of pixels with the number of PEs on a one-to-one basis. The processor handles *predictable* packets, in terms of both size and number, so the

bandwidth is also itself predictable. Also, it appears inevitable that the Childers processor is designed for, and is only capable of, processing one packet at any given time. There is no dynamic element in the operation of the Childers processor. If the Childers processor were to be used for a different source of data packets, the organization of the number of pixel/PE pairs would have to be redesigned and reconfigured. There is no "dynamic" aspect in a complete redesign.

By contrast, Applicants' specification describes a processor formed of a plurality of processing elements (PEs) adapted to receive an incoming stream of data packets, which are of unpredictable length/size and are therefore unpredictable in the number received in a given time. In the case of a "small" packet, it can be read into a single PE along with its header. However, where an incoming packet is "large", it is subdivided into smaller fragments or "chunks" of equal size and spread over as many PEs as are necessary to store it and process it. In both cases, a second packet is then read and put into the next set of PEs, until all PEs are filled. This is spelled out in the description, with particular reference to Figure 2, and is also specified in original claim 4. The product of packet size and packet rate is invariably constant (see page 7, lines 21), so there is no possibility of the incoming stream having large packets and a high packet rate at the same time.

This "bandwidth" feature plays an important part in the dynamic operation of Applicants' processor which is simply not possible in any of the prior art references under consideration. Applicants' processor operates without advance knowledge of the size of packets in the incoming data stream but nevertheless automatically subdivides a larger packet into fragments/chunks, which are then allocated to as many PEs as are necessary to store it for a given bandwidth. Typically, only the chunk with the header is operated on in the PE to which it is sent whereas the chunks not containing the packet header are simply stored in the relevant PEs. However, all the chunks, including the one containing the packet header, are usually assigned a prepend header to assist correctly "stitching" them together after processing. There is no need for the PEs that are merely storing chunks to be located anywhere in particular in the processor since the prepend header assigned to the chunks assists in their retrieval and onward transmission after the chunk with the packet header has been processed. It is possible, though, that all of the chunks can be processed while being held in their respective PEs, in dependence on either the data in the packet or information about the packet.

Applicants' distribution of whole (small) packets to a single PE or chunks of a larger packet to a plurality of PEs on a dynamic basis is therefore based on data bandwidth rather than pixel number. Residual chunks left over after subdivision of large packets are also sent to PEs, possibly padded out with zeros to fill the available storage space. It may even be the case that some PEs remain empty after this distribution exercise. These features are specified in various passages throughout the specification, such as page 12, line 16; page 13, lines 6-9; page 22, lines 19-21; page 23, lines 26-32; page 24; page 51; and page 59.

Applicants' architecture has a unique and significant advantage over the prior art in general and Childers in particular, in that the PE array can handle more than one packet at a time. The Childers processor only operates on one video line (i.e., packet) at any one time. It is an important feature of Applicants' invention that the processor can process many short packets or fewer long packets, such that the processor dynamically adapts to the bandwidth. At any given bandwidth (data rate) the number of packets to be processed is inversely proportional to their size, as is the number of processing elements. In other words, if the packets are short there are more of them, but each one is allocated to fewer PEs, so more packets can be processed at once. If the packets are long, on the other hand, then they use up more PEs, perhaps even all the PEs, but that is not problematic because they do not have to be processed as quickly.

As a consequence, Applicants' processor can afford to use some PEs simply to store chunks. Normally with parallel and especially SIMD processors, there is considerable concern with efficiency: that is, to ensure that all the processors are kept as busy as possible. In contrast, Applicants' processor is not influenced by such considerations and does not need to take them into account; all that matters is that the required processing can be done in the time available. The dynamic aspect of the processor takes care of this automatically. Applicants' processor processes packets at the speed they come through. The number of PEs is chosen to provide enough processing power to do whatever is needed with the PEs. If more (or less) processing needs to be done, the number of PEs in the design can be changed at the design stage (a design time change) or the size of the chunks can be changed. Bigger chunks mean more packets in the array at a time and therefore more time to process them.

Childers does not disclose or teach such an architecture, whether alone or in combination with Campbell. In particular, it is respectfully submitted that Childers does not teach a data processing architecture operable to process more than one packet at a time.

Childers, as has been explained above, is designed for handling a line of video at a time, one PE per pixel and each line constituting one packet. Distribution of the packet to the one-dimensional PE array is pre-determined in that each PE handles data relating to a single respective pixel of a packet. There is no scope for dynamic distribution of packets and packet chunks to the PEs. Any relationship that exists between the bandwidth of the incoming data and the numbers of PEs to which the data is distributed is pre-determined, based on the knowledge that the processor's intended use is to process lines of video data, one line at a time. There is no unpredictability in the incoming data, so the notion of providing a dynamic distribution is irrelevant and would not have occurred to one of ordinary skill in the art. There would therefore have been absolutely no incentive or motivation for the skilled person even to contemplate consulting Campbell for a solution to a problem that does not exist, according to Childers. The Office's argument that one of ordinary skill would have turned to Campbell to determine the number of PEs based on the packet size and to allocate the data based on the available number of PEs is beside the point at least because there is still no dynamic aspect to the combined teaching of Childers and Campbell, there is still no unpredictability, and there is still no ability to handle more than one packet at a time. It is therefore respectfully submitted that claim 1, as currently amended, is patentably distinguishable over Childers in view of Campbell.

The other claims rejected under this heading (i.e., claims 2-9, 11, 17, 19, 20-22, 32, 34, 38, 40, 44, 46, 50, 52, 56, 59 and 60) are variously dependent from claim 1. Accordingly, it is respectfully submitted that these claims are patentably distinguishable over any combination of Childers with Campbell for at least the reasons set forth above with respect to claim 1.

For at least the foregoing reasons, it is respectfully contended that claims 1-9, 11, 17, 19, 20-22, 32, 34, 38, 40, 44, 46, 50, 52, 56, 59, and 60 are patentable over the prior art of record. Accordingly, it is respectfully requested that the rejection of these claims under 35 U.S.C. §103(a) be withdrawn.

Claims 12, 13, 35, 41, 47, and 53 again stand rejected under 35 U.S.C. §103(a) as allegedly being unpatentable over Childers et al. in view of Gove et al. (US Patent 5,371,896) (henceforth "Gove"). This rejection is respectfully traversed.

Claim 12 depends from claim 1, and thereby incorporates all of the features of that base claim. Claim 12 is therefore patentably distinguishable over Childers for at least the

reasons set forth above. It is noted that the Office has incorporated by reference its acknowledgement, in numbered paragraph 4 of the January 9, 2006 Office Action, that "Childers did not specifically show the number of PEs was determined based on the size of the data packet as claimed."

Gove also fails to disclose an input device "operable to distribute whole data packets of potentially varying size across one or more processing elements such that the number of processing elements across which each whole data packet is distributed is dynamically determined based at least in part on the size of the whole data packet," and therefore fails to make up for the deficiencies of Childers. Because these features would be lacking in any combination of Childers with Gove, such a combination would fail to support a *prima facie* case of obviousness.

Moreover, Applicants still respectfully maintain that the Office's reliance on Gove as disclosing claim 12's recitation of a data processing architecture "wherein at least one processing element is operable to enter a standby mode of operation in dependence upon data received by that processing element" is unfounded. Applicants' claim 12 defines the processing element as being operable to enter a standby mode of operation in dependence upon data received by that processing element." (Emphasis added.) (For example, whether a given one of the processing elements is activated or idle may depend on whether that processing element contains the packet header. See, e.g., specification at page 12, lines 15-22.) As explained in Applicants' previously-filed remarks, Gove discloses a processor switchable between SIMD and MIMD using a cross-bar switch interconnecting various PEs and memory to change the combination of distributed/shared memory. There are several parallel processors interconnected with the memory banks. When there is contention (i.e., simultaneous accesses to RAM by any two system devices) a "SIMD pause" signal is routed to pause *all* PEs (see Gove Fig. 30). *The pause is therefore not dependent on data, as required by the Applicants' claim 12*, but is instead dependent on the occurrence of contention in a separate unit. The standby mode required of Applicants' claim is therefore not triggered in the way disclosed by Gove.

Claims 13, 35, 41, 47, and 53 depend from claim 12, and are therefore patentably distinguishable over the Gove patent for at least the same reasons as those set forth above.

In view of the foregoing, it is respectfully asserted that claims 12, 13, 25, 41, 47, and 53 are patentable over the prior art of record. Accordingly, it is respectfully requested that the rejection of these claims under 35 U.S.C. § 103(a) be withdrawn.

Claims 23-31 and 58 again stand rejected under 35 U.S.C. §§ 102(a) and (b) as allegedly being anticipated by Horst (US Patent 5,404,550). This rejection is respectfully traversed.

Claim 23 defines an input/output system for transferring data to and from a plurality of processing elements arranged in a single instruction multiple data (SIMD) array, the system being operable to transfer data packets of different sizes to respective ones of the processing elements in the array.

To a large extent, this claim is tied in with claim 1, in that it comprehends the same overall concept of a processor comprising an array of PEs, wherein the PEs are operable with data packets of variable size. Applicants are here claiming the aspect that the input/output system transfers data packets of different sizes to respective PEs; that is, whole packets are delivered to respective PEs, one packet per one PE. It should be appreciated that claim 23 provides a mechanism for distributing packets to PEs in a SIMD array and is not concerned with transfer of data between PEs in the array. What happens to data packets once they have been distributed to PEs across the array is immaterial to the I/O aspect specifically claimed in claim 23.

This aspect has a different slant from that in claims 1 and 10, in which a single packet is subdivided into a plurality of fragments and the fragments sent to respective PEs of a processor array. Instead, claim 23 involves an I/O system for transferring data to and from PEs in a SIMD array such that data packets of different sizes are sent to respective ones of the PEs in the array. In other words, the PEs are large enough to accept whole packets.

The Office relies on Horst as being relevant to the novelty of claim 23. Applicants have previously responded to this rejection in earlier responses to previous Office Actions, and incorporate those arguments herein by reference. However, Applicants have received no feedback as to why the Office considers the *arguments* not to be persuasive. At best, in numbered paragraph 17 of the Final Office Action, it is stated that "SIMD was most likely in Horst because it provided the background of the invention." (Emphasis added.) However, reliance on likelihood is inapposite because a fair reading of Horst shows that it does not disclose a SIMD architecture. As is well established, a document can only anticipate a claim

if it discloses *all* the claimed limitations of that claim. The Horst patent fails to satisfy this criterion because it discloses what is described as a new architecture in which processing is performed by flowing a set of tasks through a PE network. *This network is consequently incapable of operating on the basis of a Single Instruction, Multiple Data (i.e., SIMD) philosophy.* In addition, Horst does not disclose an input/output system for transferring data to and from a plurality of PEs in a SIMD array. Instead, Horst routes message blocks *between* PEs. Further, Horst describes sending transmission packets between PEs, with each such packet carrying the state of execution from the previous PE. These are substantial differences that render claim 23 novel relative to Horst, in Applicants' respectful submission.

Claims 24-31 and 58 variously depend from claim 23, and therefore incorporate all of the features of that base claim. Accordingly, for at least the foregoing reasons claims 23-31 and 58 are believed to be patentably distinguishable over the Horst patent. It is therefore respectfully requested that the rejection of these claims under 35 U.S.C. §§ 102(a) and 102(b) be withdrawn.

Claims 14, 15, 33, 36, 37, 42, 43, 48, 49, 54, and 55 again stand rejected under 35 U.S.C. §103(a) as allegedly being unpatentable over Brown (US Patent 5,872,993) in view of Childers. This rejection has been rendered moot by the cancellation of each of these claims without prejudice or disclaimer. Accordingly, it is respectfully requested that the rejection of claims 14, 15, 33, 36, 37, 42, 43, 48, 49, 54, and 55 under 35 U.S.C. §103(a) be withdrawn.

Claim 16 stands rejected under 35 U.S.C. §103(a) as allegedly being unpatentable over Childers et al. in view of Brown. This rejection is respectfully traversed.

Claim 16 depends from claim 1, and is therefore patentably distinguishable over Childers for at least the reasons set forth above. It is noted that the Office has incorporated by reference its acknowledgement, in numbered paragraph 4 of the January 9, 2006 Office Action, that "Childers did not specifically show the number of PEs was determined based on the size of the data packet as claimed."

Brown, too, fails to disclose an input device "operable to distribute whole data packets of unpredictable size across one or more of said processing elements such that the number of said processing elements across which each whole data packet is distributed is dynamically determined based at least in part on the size of the whole data packet," and therefore fails to make up for the deficiencies of Childers. Because these features would be

lacking in any combination of Childers with Brown, such a combination would fail to support a *prima facie* case of obviousness.

Moreover, one of ordinary skill in the relevant art at the time of the invention would have found neither the necessary motivation nor the expectation of success in combining Childers with Brown because of the many difficulties encountered in applying Brown's arrangement into a SIMD environment for which it was not intended.

For at least the foregoing reasons, claim 16 is believed to be patentably distinguishable over any combination of Childers with Brown. Accordingly, it is respectfully requested that the rejection of claim 16 under 35 U.S.C. §103(a) be withdrawn.

New claims 61-64 have been added without introduction of new matter. These claims depend from claim 1, and are therefore believed to be patentably distinguishable over the prior art of record for at least the same reasons as those set forth above with respect to claim 1.

The application is believed to be in condition for allowance. Prompt notice of same is respectfully requested.

Respectfully submitted,
Potomac Patent Group PLLC

Date: January 12, 2007

By: /Kenneth B. Leffler, Reg. No. 36,075/
Kenneth B. Leffler
Registration No. 36,075

P.O. Box 270
Fredericksburg, Virginia 22404
703-718-8884